7-MAY-77 20:20:20-PDT,15178;000000000000
Mail from SU-AI rcvd at 7-MAY-77 2020-PDT
Date: 7 May 1977 2013-PDT
From: GFF at SU-AI (Geoff Goodfellow)
To: Hart at SRI-KL, Sacerdoti at SRI-KL, Lynch at SRI-KL
To: Lederberg at SUMEX-AIM

"Computer"
by LEE DEMBART
(c) 1977 New York Times News Service
1:50 pm, 7 May 1977

NEW YORK - Can machines think? Should they? The computer world is in the midst of a fundamental dispute over those questions because an eminent computer scientist wrote a book arguing that machines could never be made to reason like people and, what was more, should not be.

Twenty years ago, in the infancy of the computer revolution, before problems began cropping up, the public ws told that computers would be smarter than brains. Computer chess playing and machine translation, for example, were just around the corner. So far, neither has been accomplished successfully, and neither is likely to be any time soon.

Nor have computers had much success in making decisions that require judgment. They can rattle off the Manhattan telephone directory unerringly time after time, which no human can do, but they cannot begin to distinguish one face from another, as babies can.

Computer scientists have always said, ''Give us more time, the problem is more complex than we thought.'' Then one of them Joseph Weizenbaum, a professor of computer science at the Massachusetts of Technology, wrote a book saying that the project was fundamentally unsound and dangerous to pursue, partly, he said, because the computers' and humans' ways of thought would always be alien, and because knowledge might become limited to what a computer could understand.

The elders of the artificial intelligence community reacted with outrage. Even those who agreed with his premises criticized the book as being too harsh in tone, too personal in its attacks. Computer journals have bristled over the last year with reviews, comments and replies provoked by Weizenbaum's book, ''Computer Power and Human Reason'' (W.H. Freeman & Co., 1976). Now the controversy has spilled into the prestigious publication Science, in whose pages he was attacked several weeks ago.

In an article entitled ''What computers Mean for Man and Society,'' Prof. Herbert A. Simon of the Carnegie-Mellon University in Pittsburgh argued that computers were no more or less dangerous than any other machine of the industrial revolution.

Simon, who is one of the leading figures in artificial intelligence research, asserted that Weizenbaum's position was that man was not subject to natural laws and that knowledge about man's ability to think was dangerous.

''I have pronounced heresy and I am a heretic,'' the tweedy, 54-year-old Weizenbaum said in his office in Cambridge, Mass.

Weizenbaum, who says he loves beautiful machinery, was not the first critic of artificial intelligence, but he was the first from inside the field.

''Weizenbaum has as much experience as almost anyone in the world with this subject,'' one of his opponents, Prof. Bruce G. Buchanan, acknowledged. ''He is a much more informed critic than we have ever had.''

did not understand what was being said.

Otherwise, the machine had a series of canned sentences that it could generate when it encountered words or structures that it recognized.

Dr. Ken Kolby (sic) of Stanford built a program called PARRY, which responded as a paranoid would, and the ''paranoid'' and the ''psychiatrist'' were set to talking, generating a conversation that included the following circular exchange:

PARRY: In my opinion you have a one-track mind.

ELIZA: What makes you think I have a one-track mind?

PARRY: I used to go to the track every week.

ELIZA: What does that suggest to you?

PARRY: I avoid racetracks now because the races are fixed.

ELIZA: Is that the real reason?

PARRY: I'd rather not discuss it.

ELIZA: What else would you like to discuss?

PARRY: That's something I don't want to go into.

After writing the program, Weizenbaum said, he was startled to find that instead of recognizing it as a parlor trick, some people thought they were talking to a person, and some computer scientists predicted that this tool would bring computer psychiatry to large numbers of people now unable to afford treatment.

''There are aspects to human life that a computer cannot understand - cannot,'' Weizenbaum said. ''It's necessary to be a human being. Love and loneliness have to do with the deepest consequences of our biological constitution. That kind of understanding is in principle impossible for the computer.''

But Lederberg of Stanford said he felt that computer psychiatry should be explored. ''I don't find it at all unimaginable that it may be very useful,'' he said in an interview in his ofe, which has two computer terminals (Weizenbaum's has one).

''I don't see anything either inappropriate or immoral in it given the framework that it's a help to an individual to help himself,'' Lederberg said. ''The needs for mental health support are so enormous and are so impossible to meet by using people as the sole basis that some amplification by having machines as amplifiers may be worth investigating.''

But with MYCIN, the program that diagnoses disease and prescribes treatment, the experts were not willing to do away with the doctor and let the computer dispense pills to each patient in the hospital.

''If a program such as MYCIN were acting independently of a physician, I think that would be inappropriate,'' said Buchanan, who has a computer terminal in his home. ''The problem is that human diseases are open-ended.

''Our goal is to build a program that can assist working scientists with reasoning problems. You wouldn't expect a tool to have all of the power of a working scientist, but you would hope that you could design a smart system to provide some of the solutions to subproblems.''

But the early goal of artificial intelligence studies, to build a machine that could handle unrestricted problems within a given area, has proved more elusive with each advance of knowledge. A computer like Hal in the movie ''2001,'' which talks, listens, thinks, reasons, has emotions and so forth, is nowhere within the conceivable future.

In chess, there are computers programmed to play a good game. But those which are based on artificial intelligence principles, which seek to apply strategy rather than pure memory, are not among them. The memory programs, which defeat those with artificial intelligence,

While machines provide the ostensible cause of the debate, it has encompassed some of mankind's oldest and newest questions: How do we know what we know? What does it mean for a human to know something? What does it mean for a computer? What is creativity? How do we think? What are the limits of science? What are the limits of digital computers?

There are also public policy questions, for the vast majority of work in artificial intelligence is supported by the federal government through agencies like the Defense Department and the National Aeronautics and Space Administration, with grants estimated at $5 million a year. Efforts to get computers to reason involve between 200 and 500 people in this country.

At Stanford University in Palo Alto, Calif., one of a handful of centers researching artificial intelligence, computer scientists conceded that they have not achieved the major breakthrough of teaching machines to think, rather than to simply repeat what they have been told, but they eagerly demonstrated several programs that, they said, exhibited ''cleverness.''

The ultimate aim of one such program, called MYCIN, is to enable a physician to type in a patient's symptoms and test results and get back a diagnosis and suggested treatment.

So far, the machine has been taught 450 rules regarding meningitis, culled from medical journals and Stanford hospital case histories. It compares patients' histories to those rules and has given accurate diagnoses under experimental conditions. Sometimes the machine will ask for more information about a test or about the patient, if the given information does not fulfill its program, and will explain to the physician why it wants to know. It is also able to explain how it reaches its conclusion by telling which of its rules apply to the given patient.

''These are programs with very narrow limits,'' said Buchanan, who worked on MYCIN with Dr. Joshua Lederberg, professor of genetics at Stanford and Nobel laureate in medicine. ''There's no widespread intelligence here. Yet.''

Weizenbaum retorted from MIT, ''My argument that it cannot be done does not rest on any particular limitation of computers, or the state of the art, or anything of that kind. It simply rests on the nature of science.''

''How far can science see into anything, particularly man?'' he said. ''The artificial intelligentsia claims that the computer metaphor can somehow explicate the whole man. This strikes me as being enormously arrogant and just plain wrong - wrong in principle.''

Back at Stanford, Lederberg countered, ''Assertions about what will never be possible are not logically tenable. I don't know what's impossible. I prefer to work in what is possible and try to see what we can do.''

''Organisms are built out of matter. I can see no place in principle by which I can believe that a brain can do a computation that a machine cannot,'' he said. ''A brain is, however, enormously more complex and dynamically organized, built on very different principles and subject to four billion years of evolution. So far we're far from having a very clear picture of how that is put together. That's one of the major frontiers of biology at the present time.''

One of the reasons why Weizenbaum became skeptical that parallels could be drawn between human and computer thought was the reaction to a major artificial intelligence program he wrote in 1965. Called ELIZA, it aimed to simulate a psychiatrist at work. The psychiatrist's role was chosen for this attempt at conversation because the computer could always respond with ''Please go on'' if it

do not attempt to simulate human thought. They work because of the brute force of the computer in its ability to examine many moves ahead, just as a 747 flies because of the brute force of its engines and not because it simulates a bird in flight.

In machine translation, the problem seemed simple for a computer: Give it a dictionary and a grammar of each language and let it translate. But no general purpose machine has yet been devised that can figure out a sentence like ''The book is in the pen.'' No computer can yet can read a novel and write a summary of its plot, much less translate it into another language.

Some of these issues were discussed five years ago by a philosopher at the University of California at Berkeley, Hubert L. Dreyfus, whose book ''What Computers Can't Do'' (Harper & Row, 1972) created a small stir then. Dreyfus, who freely quotes the German phenomenologist Martin Heidegger in support of his views, was dismissed by the computer people as an outsider.

In an interview recently in his spare office at Berkeley, Dreyfus argued that human activity was understandable only as a whole and could not be described as a complex assembly of simple parts. ''Man is not a mechanism,'' he said.

''If what they want are localized, restricted, game-like situations, they can make solutions,'' he said. ''But if they think they can go on broadening these little solutions till they get to the big ones, I think they just haven't faced the situation.

''Unfortunately, we cannot have a science of human beings. That's what this comes down to. Even though Plato had a dream that we could and Newton had success with physical objects.''

Lederberg, who for the last year and a half has been working on building a computer program that will help plan biological experiments, disagreed, at least in theory. He said that computers just have not been big enough to handle the complexity necessary for the big solutions. The researchers said that they have been able to set up programs that mimic human problem-solving strategies within limited contexts. ''Efforts to keep bootstrapping to higher and higher levels of generality really are impeded by the hardware,'' Lederberg said.

(In addition to hardware, which is the computer, and software, which is the program, computer scientists have lately begun talking about ''wetware,'' which is the human brain.)

Lederberg said that he thought the computer's ability to combine and examine large numbers of alternative possibilities could theoretically enable it to use human problem-solving strategies.

''I don't know how I get 'flashes of insight,' '' he said. ''I'm willing to go along with Einstein's self-description that he thinks it's combinatorial play, that he thinks it's parsing through very large numbers of combinatorial alternatives. But it's got to be structured in some way.''

The computer scientists noted that their failures in artificial intelligence have been helpful because they have helped specialists in other areas make explicit the knowledge that they have. And it has helped psychologists studying information processing in the ...

''... you can codify human thinking may not tell you that that's exactly the way humans do it, but it may give you some insight into how humans might think about the problem to do it better,'' Buchanan said.

But Weizenbaum said he believed that it was dangerous to pursue the goal of artificial intelligence, for machines will always be alien to humans, though they may eventually appear to be smarter.

He traces his views of machines to his experience as a small boy in

Nazi Germany. In 1936, at the age of 13, he fled with his parents to the United States, settling in Detroit.

''I had an introduction to the world in formative years of the miscarriage of the ultimate form of rationality,'' Weizenbaum said.

''What decided me to go into mathematics,'' he said, ''was that of all the things that one could study, mathematics seemed by far the easiest. Mathematics is a game. It is entirely abstract. Hidden behind that recognition that mathematics is the easiest is the corresponding recognition that real life is the hardest. That has been with me since childhood.''

Lederberg agreed. ''There are things that should not be done.'' But, he added, ''I don't see the difference between things that people shouldn't do and things that computers shouldn't do. What should not be done is to allow computers to get out of control.''

Said Buchanan: ''Any tool with sufficient power to have enormous benefits also carries elements of risk if that tool is misused. Scientists who are developing that tool need to be aware of the risks as well as of the benefits.''